

# Verifying time, memory and communication bounds in systems of reasoning agents

Natasha Alechina   Brian Logan   Nguyen Hoang Nga   Abdur Rakib

School of Computer Science  
University of Nottingham  
Nottingham NG8 1BB, UK  
{nza,bsl,hnn,rza}@cs.nott.ac.uk

## ABSTRACT

We present a framework for verifying systems composed of heterogeneous reasoning agents, in which each agent may have differing knowledge and inferential capabilities, and where the resources each agent is prepared to commit to a goal (time, memory and communication bandwidth) are bounded. The framework allows us to investigate, for example, whether a goal can be achieved if a particular agent, perhaps possessing key information or inferential capabilities, is unable (or unwilling) to contribute more than a given portion of its available computational resources or bandwidth to the problem. We present a novel temporal epistemic logic, *BMCL*, which allows us to describe a set of reasoning agents with bounds on time, memory and the number of messages they can exchange. The bounds on memory and communication are expressed as axioms in the logic. As an example, we show how to axiomatize a system of agents which reason using resolution and prove that the resulting logic is sound and complete. We then show how to encode a simple system of reasoning agents specified in *BMCL* in the description language of a model checker, and verify that the agents can achieve a goal only if they are prepared to commit certain time, memory and communication resources.

## Categories and Subject Descriptors

I.2 [Artificial Intelligence]; F.3 [Logics and Meanings of Programs]

## General Terms

Theory

## Keywords

Formalisms and logics

## 1. INTRODUCTION

A key application of multiagent systems research is distributed problem solving. Distributed approaches to problem solving allow groups of agents to collaborate to solve problems which no single agent could solve alone (e.g., because no single agent has all the information necessary to solve the problem), and/or to solve

**Cite as:** Verifying time, memory and communication bounds in systems of reasoning agents, Natasha Alechina, Brian Logan, Nguyen Hoang Nga and Abdur Rakib, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 736-743. Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

problems more effectively, e.g., in less time than a single agent. For a given problem and system of reasoning agents, many different solution strategies may be possible, each involving different commitments of computational resources (time and memory) and communication by each agent. For different multiagent systems, different solution strategies will be preferred depending on the relative costs of computational and communication resources for each agent. These tradeoffs may be different for different agents (e.g., reflecting their computational capabilities or network connection) and may reflect the agent's commitment to a particular problem. For a given system of agents with specified inferential abilities and resource bounds it may not be clear whether a particular problem can be solved at all, or, if it can, what computational and communication resources must be devoted to its solution by each agent. For example, we may wish to know whether a goal can be achieved if a particular agent, perhaps possessing key information or inferential capabilities, is unable (or unwilling) to contribute more than a given portion of its available computational resources or bandwidth to the problem.

There has been considerable work in the agent literature on distributed problem solving in general (for example, [18, 23, 25, 28]) and on distributed reasoning in particular ([1, 11]). Much of this work analyses the time and communication complexity of distributed reasoning algorithms. However, while we have upper (and some lower) bounds on time and memory requirements for reasoning in distributed systems, we lack tools for reasoning about tradeoffs between computational and communication resources. In this paper we present a framework for reasoning about tradeoffs between time, memory and communication in systems of distributed reasoning agents. We introduce a novel epistemic logic, *BMCL*, for specifying resource-bounded reasoners. Critically, the logic allows upper bounds on the resource commitments (time, memory and communication) of each agent in the system to be specified. We prove that the logic is sound and complete. Using simple resolution examples, we show how to encode systems of distributed reasoning agents specified in the logic in a model checker, and verify some example properties. In contrast to previous work, e.g., [4, 5, 6, 2]) which focused primarily on memory limitations of single reasoners, our approach allows us to specify bounds on the number of messages the agents can exchange, allowing the investigation of tradeoffs between different resources. This allows us to determine whether, for example, giving a reasoner more memory will result in a shorter proof, or whether communication between agents can reduce either time or memory requirements or both.

The structure of the paper is as follows. In section 2 we introduce the problem of resource bounds in distributed reasoning and in section 3 we explain how we measure time, space and communication

costs for a distributed reasoning problem. In section 4 we introduce the epistemic logic *BMCL*. Model-checking experiments are described in section 5. We survey related work in section 6 and conclude in section 7.

## 2. DISTRIBUTED REASONERS

We define the ‘shape’ of a proof in terms of the maximum space requirement at any step in the proof and the number of inference steps it contains. The lower bound on space for a given problem is then the least maximum space requirement of any proof, and the lower bound on time is the least number of inference steps of any proof. In general, a minimum space proof and a minimum time proof will be different (have different shapes). Bounding the space available for a proof will typically increase the number of inference steps required and bounding the number of steps will increase the space required. For example, a proof which requires only the minimum amount of space may require rederivation of intermediate results.

We define the bounds on a reasoning agent in terms of its available resources expressed in terms of memory, time and communication. We assume that the memory required for a particular proof can be taken to be its space requirement (e.g., the number of formulas that must be simultaneously held in memory) times some constant, and the number of inference steps executed times some constant can be taken as a measure of the time necessary to solve the problem. The communication requirement of a proof is taken to be the number of messages exchanged with other agents. In what follows, we ignore the constants and assume that the units of problem size and resources are the same.

For a particular agent solving a particular problem, the space available for any given proof is ultimately bounded by the size of the agent’s memory and the number of inference steps is bounded by the time available to the agent, e.g., by a response time guarantee offered by the agent, or simply the point in time at which the solution to the problem becomes irrelevant. The question then arises of whether a proof can be found which falls within the resource envelope defined by the agent’s resource bounds.

For a single agent which processes a single goal at a time, the lower bounds on space for the goal determines the minimum amount of memory the agent must have if it is to solve the problem (given unlimited time); and the lower bound on time determines the time the agent must commit to solving the problem (given unlimited memory). In the general case in which the agent is attending to multiple goals simultaneously, the memory and time bounds may be given not by the environment, but by the need to share the available resources between multiple tasks. For example, the agent may need to share memory between multiple concurrent tasks and/or devote no more than a given proportion of CPU to a given task. In both cases, the agent designer may be interested in tradeoffs between resource bounds; for example, whether more timely responses can be provided by pursuing fewer tasks in parallel (thereby making more memory available to each task) or whether more tasks can be pursued in parallel if each task is allowed to take longer.

In the distributed setting we distinguish between *symmetric problem distributions*, where all agents have the same premises, and *asymmetric problem distributions* where different premises may be assigned to different agents. We also distinguish between *homogeneous reasoners* (when all agents have the same rules of inference and resource bounds) and *heterogeneous reasoners*, (when different agents have different rules of inference and/or resource bounds).

Distribution does not necessarily change the shape (maximum space requirement and number of inference steps) of a proof. How-

ever, in a distributed setting the tradeoffs between memory and time bounds are complicated by communication. Unlike memory and time, communication has no direct counterpart in the proof. However like memory, communication can be substituted for time (e.g., if part of the proof is carried out by another agent), and, like time, it can be substituted for memory (e.g., if a lemma is communicated by another agent rather than having to be remembered). In the distributed setting, each agent has a minimum memory bound which is determined by its inference rules and which may be smaller than the minimum space requirement for the problem. If the memory bound for all agents taken individually is less than the minimum space requirement for the problem, then the communication bound must be greater than zero. (If the memory bound for all agents taken together is less than the minimum space requirement for the problem, then the problem is insoluble for any communication bound).

With a symmetric problem distribution, if the memory bound for at least one agent is greater than the minimum space requirement for the problem, the minimum communication bound is zero (with unbounded time). If the problem distribution is asymmetric, i.e., not all agents have all the premises, then the lower bound on communication may again be non-zero, if a necessary inference step requires premises from more than one agent.

In the next section, we present measures of space, time and communication for distributed reasoning agents which allow us to make these tradeoffs precise.

## 3. MEASURING RESOURCES

We assume a set of  $n$  agents. Each agent  $i$  has a set of propositional inference rules  $R_i$  (for example,  $R_i$  could contain conjunction introduction and modus ponens, or it could contain just a single rule of resolution) and a set of premises  $K_i$ .

For a single agent, the notion of a derivation, or a proof of a formula  $G$  from  $K_i$  is standard, and the time and space complexity of proofs are well studied. Our model of space complexity is based on [9]. We view the process of producing a proof of  $G$  from  $K_i$  as a sequence of *configurations* or states of a reasoner, starting from an empty configuration, and producing the next configuration by one of the following operations:

**Read** copies a formula from  $K_i$  into the current configuration (possibly overwriting a formula from the previous configuration)

**Infer** applies a rule from  $R_i$  to formulas in the current configuration (possibly overwriting a formula from the previous configuration)

The sequence of configurations constitutes a proof of  $G$  if  $G$  appears in the last configuration. Time complexity corresponds to the length of the sequence, and space complexity to the size of configurations.<sup>1</sup> The size of a configuration can be measured either in terms of the number of formulas appearing in the configuration or in terms of the number of symbols required to represent the configuration. Clearly, for some inference systems, for example, where the set of inference rules contains both conjunction introduction and conjunction elimination, the first size measure results in constant space usage. However, for other systems, such as resolution, counting formulas results in non-trivial space complexity [16]. In this paper, we take the size of a configuration to be the maximal

<sup>1</sup>Note that we deviate from [9] in that we do not have an explicit *erase* operation, preferring to incorporate erasing (overwriting) in the *read* and *infer* operations. This obviously results in shorter proofs; however including an explicit erase operation gives proofs which are no more than twice as long as our proofs if we don’t require the last configuration to contain *only* the goal formula.

number of formulas, since all the reasoning systems we consider have a non-trivial space complexity for this measure.

#	Configuration	Operation
1	{ }	
2	{ $A_1$ }	<b>Read</b>
3	{ $A_1, A_2$ }	<b>Read</b>
4	{ $A_1, A_1 \wedge A_2$ }	<b>Infer</b>
5	{ $A_1 \wedge A_2, A_1 \wedge A_2 \rightarrow B_1$ }	<b>Read</b>
6	{ $A_1 \wedge A_2, B_1$ }	<b>Infer</b>

**Figure 1: Example derivation using  $\wedge_I$  and  $MP$**

As an illustration, Figure 1 shows the space and time complexity of the derivation of the formula  $B_1$  from  $A_1, A_2, A_1 \wedge A_2 \rightarrow B_1$  in an inference system which contains only conjunction introduction and modus ponens. The length of the proof is 6 and the space usage is 2 (at most 2 formulas need to be present in the configuration at any given time). It is worth observing that the inference system consisting of just conjunction introduction and modus ponens does not have constant space complexity when space is measured as the number of formulas; a sequence of derivation examples requiring (logarithmically) growing space can easily be constructed starting from the example above, and continuing with a derivation of  $C_1$  from  $A_1, A_2, A_3, A_4, A_1 \wedge A_2 \rightarrow B_1, A_3 \wedge A_4 \rightarrow B_2, B_1 \wedge B_2 \rightarrow C_1$ , etc.

#	Configuration	Operation
1	{ }	
2	{ $A_1 \vee A_2$ }	<b>Read</b>
3	{ $A_1 \vee A_2, \neg A_1 \vee A_2$ }	<b>Read</b>
4	{ $A_1 \vee A_2, A_2$ }	<b>Infer</b>
5	{ $A_2, A_1 \vee \neg A_2$ }	<b>Read</b>
6	{ $A_2, A_1 \vee \neg A_2, \neg A_1 \vee \neg A_2$ }	<b>Read</b>
7	{ $A_2, \neg A_2, \neg A_1 \vee \neg A_2$ }	<b>Infer</b>
8	{ $\emptyset, \neg A_2, \neg A_1 \vee \neg A_2$ }	<b>Infer</b>

**Figure 2: Example derivation using resolution**

Most research in time and space complexity of proofs has focused on the lower bounds for the inference system as a whole. While we are interested in the lower bounds, we are also interested in the trade-offs between time and space usage for particular derivations. For example, consider a set of premises  $A_1, A_2, A_3, A_4, A_1 \wedge A_2 \rightarrow B_1, A_3 \wedge A_4 \rightarrow B_2, B_1 \wedge B_2 \rightarrow C_1$  and a goal formula  $A_1 \wedge A_2 \wedge C$ . It is possible to derive the goal from the premises using conjunction introduction and modus ponens and configurations of size 3 in 17 steps (deriving  $A_1 \wedge A_2$  twice). On the other hand, with configurations of size 4 the proof is 3 steps shorter.

Different inference systems have different complexity and different tradeoffs. Figure 2 illustrates the (non-trivial) space complexity of resolution proofs in terms of the number of formulas in a configuration. The example, which is due to [16], shows the derivation of an empty clause by resolution from the set of all possible clauses of the form

$$\sim A_1 \vee \sim A_2 \vee \dots \vee \sim A_n$$

(where  $\sim A_i$  is either  $A_i$  or  $\neg A_i$ ), for  $n = 2$ . Its space usage is 3 and the length of the proof is 8.

In the multiagent case, when several reasoners can communicate to derive a common goal, an additional resource of interest is how many messages the reasoners must exchange in order to derive the goal. In the distributed setting, we assume that each agent has its own set of premises and inference rules and its own configuration, and that the reasoning of the agents proceeds in lock step. In

addition to **Read** and **Infer**, each reasoner can perform two extra operations:

**Skip** which leaves its configuration unchanged

**Copy** if agent  $i$  has a formula  $A$  in its current configuration, then agent  $j$  can copy it to its next configuration

The goal formula is derived if it occurs in the configuration of one of the agents. Our model of communication complexity is based on [29], except that we count the number of formulas exchanged by the agents rather than the number of bits exchanged. The communication complexity of a joint derivation is then the (total) number of **Copy** operations in the derivation.

#	Agent 1		Agent 2	
	Configuration	Op.	Configuration	Op.
1	{ }		{ }	
2	{ $A_1 \vee A_2$ }	<b>Read</b>	{ $A_1 \vee \neg A_2$ }	<b>Read</b>
3	{ $A_1 \vee A_2, \neg A_1 \vee A_2$ }	<b>Read</b>	{ $\neg A_1 \vee \neg A_2, A_1 \vee \neg A_2$ }	<b>Read</b>
4	{ $A_1 \vee A_2, A_2$ }	<b>Infer</b>	{ $\neg A_2, A_1 \vee \neg A_2$ }	<b>Infer</b>
5	{ $A_1 \vee \neg A_2, A_2$ }	<b>Read</b>	{ $\neg A_2, A_2$ }	<b>Copy</b>
6	{ $A_1, A_2$ }	<b>Infer</b>	{ $\{\}, A_2$ }	<b>Infer</b>

**Figure 3: Example derivation using resolution with two agents**

In general, in a distributed setting, trade-offs are possible between the number of messages exchanged and the space (size of a single agent's configuration) and time required for a derivation. The total space use (the total number of formulas in all agent's configurations) clearly cannot be less than the minimal configuration size required by a single reasoner to derive the goal formula from the union of all knowledge bases using all of the available inference rules, however this can be distributed between the agents in different ways, resulting in different numbers of exchanged messages. Similarly, if parts of a derivation can be performed in parallel, the total derivation will be shorter, though communication of the partial results will increase the communication complexity. As an illustration, figure 3 shows one possible distribution of the resolution example in figure 2. As can be seen, two communicating agents can solve the problem faster than a single agent.

## 4. A BOUNDED MEMORY AND COMMUNICATION LOGIC $BMCL$

In this section we present  $BMCL$ , a temporal epistemic logic which allows us to describe a set of reasoning agents with bounds on memory and on the number of messages they can exchange. In this logic, we can express statements like ‘the agents will be able to derive the goal formula in  $n$  inference steps’. The bounds on memory and communication ability are expressed as axioms in the logic. In this paper, as an example, we have chosen to axiomatise a set of agents reasoning using resolution. Other reasoning systems can be axiomatised in a similar way, and we briefly sketch how to add model conditions and axioms for reasoners which reason using conjunction introduction and modus ponens to our logic at the end of this section.

Let the set of agents be  $AG = \{1, 2, \dots, n_{AG}\}$ . For simplicity, we assume that they agree on a finite set  $PROP$  of propositional variables (this assumption can easily be relaxed, so that only some propositional variables are shared). Since each agent uses resolution for reasoning, we assume that all formulas of the internal language of the agents are in the form of *clauses*. For convenience, we define a clause as a set of *literals* in which a literal is a propositional variable or its negation. Then the set of literals is defined as  $LPROP = \{p, \neg p | p \in PROP\}$ . If  $L$  is a literal, then

$\neg L$  is  $\neg p$  if  $L$  is a propositional variable  $p$ , and  $p$  if  $L$  is of the form  $\neg p$ . Let  $\Omega$  be the set of all possible clauses over  $PROP$ , i.e.,  $\Omega = \wp(LPROP)$ . Note that  $\Omega$  is finite.

The only rule of inference that each agent has is the *resolution rule* which is defined as follows:

$$\frac{\alpha \ni L \quad \beta \ni \neg L}{(\alpha \setminus \{L\}) \cup (\beta \setminus \{\neg L\})} \text{Res}$$

which states that if there are two clauses  $\alpha$  and  $\beta$  such that one contains a literal  $L$  and the other contains  $\neg L$ , then we can derive a new clause  $(\alpha \setminus \{L\}) \cup (\beta \setminus \{\neg L\})$ .

Each agent  $i$  has a memory of size  $n_M(i)$  where one unit of memory corresponds to the ability to store an arbitrary clause. Agent  $i$  can read clauses from its set of premises  $K_i$ . We assume that each  $K_i$  is finite. The communication ability of the agents is expressed by the *copy* action which copies a clause from another agent's memory. The limit on each agent's communication ability is  $n_C(i)$ : in any valid run of the system, agent  $i$  can perform at most  $n_C(i)$  copy actions.

## 4.1 Syntax of $BMCL$

The syntax of  $BMCL$  is defined inductively as follows.

- $\top$  is a well-formed formula (wff) of  $BMCL$ .
- If  $\alpha$  is a clause, then  $B_i^r \alpha$  is a wff of  $BMCL$ , for all  $i \in AG$ .
- If  $\alpha$  is a clause, then  $B_i^c \alpha$  is a wff of  $BMCL$ , for all  $i \in AG$ .
- If  $\phi$  and  $\psi$  are wff, then so are  $\neg \phi$ ,  $\phi \wedge \psi$ .
- If  $\phi$  and  $\psi$  are wff, then so are  $X\phi$ ,  $\phi U \psi$ ,  $Y\phi$ ,  $\phi S \psi$  and  $A\phi$ .

Classical abbreviations for  $\vee$ ,  $\rightarrow$ ,  $\leftrightarrow$  and  $\perp$  are defined as usual.

The language contains both temporal and epistemic modalities. For the temporal part of  $BMCL$ , we have  $PCTL^*$ , a branching time temporal logic with the past operator.<sup>2</sup> Intuitively,  $PCTL^*$  describes infinite trees, or all possible runs of the system, over discrete time. In the temporal logic part of the language,  $X$  stands for next step,  $U$  for until,  $Y$  for previous step,  $S$  for since and  $A$  for 'on all paths'. We will also use abbreviations  $F\phi \equiv \top U \phi$  for 'some time in the future',  $P\phi \equiv \top S \phi$  for 'some time in the past',  $E\phi \equiv \neg A \neg \phi$  for 'on some path' and  $start \equiv \neg Y \top$  for the starting state of the system. The epistemic part of the language consists of belief modalities  $B_i^r \alpha$ , which means that agent  $i$  has read  $\alpha$  from its knowledge base or derived it, and  $B_i^c \alpha$ , which means that  $i$  has copied  $\alpha$  from another agent. We define  $B_i \alpha$  (agent  $i$  believes  $\alpha$ ) to be  $B_i^r \alpha \vee B_i^c \alpha$ .

## 4.2 Semantics of $BMCL$

The semantics of  $BMCL$  is defined by  $BMCL$  transition systems. A  $BMCL$  transition system  $M = (S, R, V^r, V^c)$  is defined as follows.

- $S$  is a non-empty set of states.
- $R \subseteq S \times S$  is a total binary relation, i.e. for any  $s \in S$ , there exists  $t \in S$  such that  $(s, t) \in R$ . Moreover, it is also required that  $(S, R)$  is a tree-frame. A branch  $\sigma$  is an infinite sequence  $(s_0, s_1, \dots)$  such that  $(s_i, s_{i+1}) \in R$  for all  $i \geq 0$ ,  $\sigma_i$  denotes the element  $s_i$  of  $\sigma$  and  $\sigma_{\leq i}$  is the prefix  $(s_0, s_1, \dots, s_i)$  of  $\sigma$ . The set of all branches is denoted as  $BR$ . Note that since  $(S, R)$  is a tree-frame every state  $s$  has a unique past  $past(s) = \sigma_{\leq i}$  where  $\sigma_i = s$ .

<sup>2</sup>The reason we use  $PCTL^*$  rather than  $CTL^*$  is that we need the past operator to express the bound on agent communication.

- $V^r : S \times AG \rightarrow \wp(\Omega)$ , is a mapping that defines for each state which formulas an agent believes due to reading or inference.
- $V^c : S \times AG \rightarrow \wp(\Omega)$ , is a mapping that defines for each state which formulas an agent copied from the memories of other agents.

The truth of a  $BMCL$  formula in a state at point  $n$  of a path  $\sigma$  of  $M$  is defined inductively as follows:

- $M, \sigma, n \models B_i^r \alpha$  iff  $\alpha \in V^r(\sigma_n, i)$ ,
- $M, \sigma, n \models B_i^c \alpha$  iff  $\alpha \in V^c(\sigma_n, i)$ ,
- $M, \sigma, n \models \neg \phi$  iff  $M, \sigma, n \not\models \phi$ ,
- $M, \sigma, n \models \phi \wedge \psi$  iff  $M, \sigma, n \models \phi$  and  $M, \sigma, n \models \psi$ ,
- $M, \sigma, n \models X\phi$  iff  $M, \sigma, n+1 \models \phi$ ,
- $M, \sigma, n \models \phi U \psi$  iff  $\exists m \geq n$  such that  $\forall k \in [n, m)$   $M, \sigma, m \models \phi$  and  $M, \sigma, m \models \psi$ ,
- $M, \sigma, n \models Y\phi$  iff  $n > 0$  and  $M, \sigma, n-1 \models \phi$ ,
- $M, \sigma, n \models \phi S \psi$  iff  $\exists m \leq n$  such that  $\forall k \in (m, n]$   $M, \sigma, m \models \phi$  and  $M, \sigma, m \models \psi$ ,
- $M, \sigma, n \models A\phi$  iff  $\forall \sigma' \in BR$  such that  $\sigma'_{\leq n} = \sigma_{\leq n}$ ,  $M, \sigma', n \models \phi$ .

Now we describe conditions on the models. The first set of conditions refers to the accessibility relation  $R$ . The intuition behind the conditions is that  $R$  corresponds to the agents executing actions  $\langle a_1, \dots, a_{n_{AG}} \rangle$  in parallel, where action  $a_i$  is a possible action (transition) for the agent  $i$  in a given state. Actions of each agent  $i$  are:  $Read_{i,\alpha,\beta}$  (reading a clause  $\alpha$  from the knowledge base and erasing  $\beta$ ),  $Res_{i,\alpha_1,\alpha_2,L,\beta}$  (resolving  $\alpha_1$  and  $\alpha_2$  on  $L$  and erasing  $\beta$ ),  $Copy_{i,\alpha,\beta}$  (copying  $\alpha$  from another agent and erasing  $\beta$ ),  $Erase_{i,\alpha}$  (erasing  $\alpha$ ), and  $Null_i$  (doing nothing), where  $\alpha, \alpha_1, \alpha_2, \beta \in \Omega$  and  $L \in LPROP$ .<sup>3</sup> Intuitively,  $\beta$  is an arbitrary clause which gets overwritten if it is in the agent's memory. If the agent's memory is full ( $|V^r(s, i)| + |V^c(s, i)| = n_M(i)$ ), then  $\beta$  has to be in  $V^r(s, i) \cup V^c(s, i)$ , otherwise we cannot add an extra formula to it (this would violate the condition on memory defined below). Not all actions are possible in any given state (for example, to perform a resolution step from state  $s$ , the agent has to have two resolvable clauses in  $s$ ). Let us denote the set of all possible actions by agent  $i$  in state  $s$  by  $R_i(s)$ .

Below is the definition of  $R_i(s)$ :

**DEFINITION 1 (AVAILABLE ACTIONS).** *For every state  $s$  and agent  $i$ ,*

1.  $Read_{i,\alpha,\beta} \in R_i(s)$  iff  $\alpha \in K_i$  and  $\beta \in \Omega$ , or if  $|V^r(s, i)| + |V^c(s, i)| = n_M(i)$  then  $\beta \in V^r(s, i) \cup V^c(s, i)$ .
2.  $Res_{i,\alpha_1,\alpha_2,L,\beta} \in R_i(s)$  iff  $\alpha_1, \alpha_2 \in \Omega$ ,  $\alpha_1 \ni L$ ,  $\alpha_2 \ni \neg L$ ,  $\alpha_1, \alpha_2 \in V^r(s, i) \cup V^c(s, i)$ ,  $\alpha = (\alpha_1 \setminus \{L\}) \cup (\alpha_2 \setminus \{\neg L\}) \notin K_i$  and  $\beta$  is as before.

<sup>3</sup>The  $Erase_{i,\alpha}$  action is introduced for purely technical reasons, to obtain a simpler axiomatisation of the system. The optimal sequences of actions found by the system when verifying properties of agents will contain no  $Erase$  actions so will not affect the verification process.

3.  $Copy_{i,\alpha,\beta} \in R_i(s)$  iff there exists  $j \neq i$  such that  $\alpha \in V^r(s, j) \cup V^c(s, j)$  and  $past(s)$  does not contain more than  $n_C(i) - 1$  transitions of the form  $Copy_{i,\beta}$ , and  $\beta$  is as before.<sup>4</sup>

4.  $Null_i$  is always in  $R_i(s)$ .

5. There are no conditions on  $Erase_{i,\alpha} \in R_i(s)$ .

Now we define effects of actions on the agent's state (assignments  $V^r(s, i)$  and  $V^c(s, i)$ ).

**DEFINITION 2 (EFFECTS OF ACTIONS).** For each  $i \in AG$ , the result of performing an action  $a$  in state  $s$  is defined if  $a \in R_i(s)$  and has the following effect on the assignment of clauses to  $i$  in the successor state  $t$ :

1. if  $a$  is  $Read_{i,\alpha,\beta}$ :  $V^r(t, i) = (V^r(s, i) \setminus \{\beta\}) \cup \{\alpha\}$  and  $V^c(t, i) = V^c(s, i) \setminus \{\beta\}$ .
2. if  $a$  is  $Res_{i,\alpha_1,\alpha_2,L,\beta}$ :  $V^r(t, i) = (V^r(s, i) \setminus \{\beta\}) \cup \{\alpha\}$  and  $V^c(t, i) = V^c(s, i) \setminus \{\beta\}$ , where  $\alpha = (\alpha_1 \setminus \{L\}) \cup (\alpha_2 \setminus \{\neg L\})$ .
3. if  $a$  is  $Copy_{i,\alpha,\beta}$ :  $V^c(t, i) = (V^c(s, i) \setminus \{\beta\}) \cup \{\alpha\}$  and  $V^r(t, i) = V^r(s, i) \setminus \{\beta\}$ .
4. if  $a$  is  $Null_i$ :  $V^r(t, i) = V^r(s, i)$  and  $V^c(t, i) = V^c(s, i)$ .
5. if  $a$  is  $Erase_{i,\alpha}$  then  $V^r(t, i) = V^r(s, i) \setminus \{\alpha\}$  and  $V^c(t, i) = V^c(s, i) \setminus \{\alpha\}$ , where  $\alpha \in \Omega$ .

**DEFINITION 3.**  $BMCM(K_1, \dots, K_{n_{AG}}, n_M, n_C)$  is the set of models  $M = (S, R, V, C)$  such that:

1. For every  $s$  and  $t$ ,  $R(s, t)$  iff for some tuple of actions  $\langle a_1, \dots, a_{n_{AG}} \rangle$ ,  $a_i \in R_i(s)$  and the assignment in  $t$  satisfies the effects of  $a_i$  for every  $i$  in  $\{1, \dots, n_{AG}\}$ .
2. For every  $s$  and a tuple of actions  $\langle a_1, \dots, a_{n_{AG}} \rangle$ , if  $a_i \in R_i(s)$  for every  $i$  in  $\{1, \dots, n_{AG}\}$ , then there exists  $t \in S$  such that  $R(s, t)$  and  $t$  satisfies the effects of  $a_i$  for every  $i$  in  $\{1, \dots, n_{AG}\}$ .
3. The bound on each agent's memory is set by the following constraint on the mappings  $V^r$  and  $V^c$ :

$$|V^r(s, i)| + |V^c(s, i)| \leq n_M(i) \text{ for all } s \in S \text{ and } i \in AG$$

Note that the bound  $n_C(i)$  on each agent  $i$ 's communication ability (no branch contains more than  $n_C(i)$   $Copy$  actions by agent  $i$ ) follows from the fact that  $Copy_i$  is only enabled if  $i$  has performed fewer than  $n_C(i)$  copy actions in the past.

### 4.3 Axiomatisation of $BMCL$

Before we give an axiomatisation for the set of models defined above, we need the following abbreviations for expressing that  $i$  has performed at least  $k$  copy actions in the past. A successful copying of a clause  $\alpha$  by agent  $i$  from an agent  $j$  is defined by the following formula:

$$copied(i, j, \alpha) \equiv B_j \alpha \wedge \neg B_i \alpha \wedge X B_i^c \alpha$$

<sup>4</sup>Assume that the state contains a communication counter for each agent  $i$ , which is set to 0 in the start state and is incremented every time  $i$  performs a copy action. After the counter reaches  $n_C(i)$ , agent  $i$  cannot perform any more copy actions.

Copying of any clause from any agent by agent  $i$  is defined as follows:

$$copied_i \equiv \bigvee_{j \in AG, \alpha \in \Omega} copied(i, j, \alpha)$$

So to say that there are at least  $k$  copy actions in agent  $i$ 's past, we can use

$$C_i^{\geq}(k) = \underbrace{(YP(copied_i \wedge YP(copied_i \wedge \dots YP(copied_i) \dots))}_{k \text{ times}}$$

and to say that there are fewer than  $k$  copy actions in agent  $i$ 's past, we can say

$$C_i^{\leq}(k) = \neg \underbrace{(YP(copied_i \wedge YP(copied_i \wedge \dots YP(copied_i) \dots))}_{k+1 \text{ times}}$$

To define exactly  $k$  copies, we can use  $C_i(0) = C_i^{\leq}(0)$  and  $C_i(k) = C_i^{\geq}(k) \wedge C_i^{\leq}(k)$  for  $k > 0$ .

Consider the following set of axiom schemata:

**A1** Axioms and rules of  $PCTL^*$  as given in [27].

**A2**  $\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \rightarrow EX(\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \wedge B_i^r \alpha)$  for all  $\alpha \in K_i$ ,  $i \in AG$ ,  $Q \subseteq \Omega$  with  $|Q| < n_M(i)$ , and  $n \geq 0$ .

**A3**  $\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \wedge B_i \alpha_1 \wedge B_i \alpha_2 \rightarrow EX(\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \wedge B_i \alpha)$  for any  $\alpha_1$  and  $\alpha_2$  such that  $\alpha_1 \ni L$  and  $\alpha_2 \ni \neg L$  for some literal  $L$ ,  $\alpha = (\alpha_1 \setminus \{L\}) \cup (\alpha_2 \setminus \{\neg L\}) \notin K_i$ ,  $Q \subseteq \Omega$  with  $|Q| < n_M(i)$ , and  $n \geq 0$ .

**A4**  $\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \wedge B_j \alpha \wedge C_i^{\leq}(n_C(i)) \rightarrow EX(\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n+1) \wedge B_i^c \alpha)$  for all  $i \neq j \in AG$ ,  $Q \subseteq \Omega$  with  $|Q| < n_M(i)$ , and  $n \geq 0$ .

**A5**  $EX(B_i \alpha_1 \wedge B_i \alpha_2) \rightarrow B_i \alpha_1 \vee B_i \alpha_2$

**A6**  $EX(\neg B_i \alpha_1 \wedge \neg B_i \alpha_2) \rightarrow (\neg B_i \alpha_1 \vee \neg B_i \alpha_2)$

**A7**  $EX(B_i^r \alpha \wedge C_i(n)) \rightarrow B_i^r \alpha \vee (\neg B_i^r \alpha \wedge C_i(n))$  for all  $\alpha \in K_i$

**A8**  $EX(B_i^r \alpha \wedge C_i(n)) \rightarrow B_i^r \alpha \vee (\neg B_i^r \alpha \wedge \bigvee_{(\alpha_1, \alpha_2) \in Res(\alpha)} (B_i \alpha_1 \wedge B_i \alpha_2 \wedge C_i(n)))$  for all  $\alpha \notin K_i$  and  $Res(\alpha) = \{(\alpha_1, \alpha_2) \in \Omega \times \Omega \mid \alpha_1 \ni L, \alpha_2 \ni \neg L \text{ and } \alpha = (\alpha_1 \setminus \{L\}) \cup (\alpha_2 \setminus \{\neg L\})\}$  for some literal  $L$  and  $n \geq 0$

**A9**  $EX(B_i^c \alpha \wedge C_i(n)) \rightarrow B_i^c \alpha \vee (\neg B_i^c \alpha \wedge C_i(n-1) \wedge (\bigvee_{j \in AG} B_j \alpha))$

**A10**  $B_i \alpha_1 \wedge \dots \wedge B_i \alpha_{n_M} \rightarrow \neg B_i \alpha_{n_M+1}$  where  $i = 1, \dots, n_{AG}$ , and  $\alpha_i \neq \alpha_j$  for all  $i \neq j$

**A11**  $C_i^{\leq}(n_C(i))$

**A12**  $\bigwedge_{j \in J} EX(\bigwedge_{q \in Q} B_j \alpha_q \wedge C_j(k_j)) \rightarrow EX \bigwedge_{j \in J} (\bigwedge_{q \in Q} B_j \alpha_q \wedge C_j(k_j))$  where  $J \subseteq AG$  and all indices  $j$  are distinct, and  $Q \subseteq \Omega$ .

**A13**  $\phi \rightarrow EX \phi$

Let  $BMCL(K_1, \dots, K_{n_{AG}}, n_M, n_C)$  be the logic defined by the our axiomatization. Then we have the following result.

**THEOREM 1.**  $BMCL(K_1, \dots, K_{n_{AG}}, n_M, n_C)$  is sound and weakly complete with respect to  $BMCM(K_1, \dots, K_{n_{AG}}, n_M, n_C)$ .

PROOF. The proof of soundness is standard. Due to lack of space, we only prove validity of the first *BMCL* axiom.

Let us consider **A2** and a model  $M = (S, R, V^r, V^c)$  of *BMCM*  $(K_1, \dots, K_{n_{AG}}, n_M)$ . Let  $\sigma = (s_0, s_1, \dots) \in BR$ . We need to prove for any  $m$ , that if  $M, \sigma, m \models \bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n)$ , then  $M, \sigma, m \models EX(\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \wedge B^r \alpha)$  where  $\alpha \in K_i$ ,  $i \in AG$ ,  $Q \subseteq \Omega$  with  $|Q| < n_M(i)$  and  $n \geq 0$ . Since  $\alpha \in K_i$ ,  $Read_{i,\alpha,\beta} \in R_i(\sigma_n)$  for some  $\beta \in \Omega \setminus Q$ . Therefore, there exists  $t \in S$  such that  $R(s, t)$  and  $t$  satisfies the effects of  $Read_{i,\alpha,\beta}$ . In other words, we obtain  $V^r(t, i) = V^r(s, i) \cup \{\alpha\} \setminus \{\beta\}$  and  $V^c(t, i) = V^c(s, i) \setminus \{\beta\}$ , this shows  $V^r(t, i) \ni \alpha$ . Since  $M, \sigma, m \models \bigwedge_{\alpha_q \in Q} B_i \alpha_q$ ,  $V^r(s, i) \cup V^c(s, i) \ni \alpha_q$  for all  $q \in Q$ . Moreover, since  $|Q| < n_M(i)$ , we have  $\beta \in \Omega \setminus Q$ , therefore  $V^r(t, i) \cup V^c(t, i) \ni \alpha_q$ . Then,  $M, \sigma', m+1 \models \bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge B^r \alpha$  for some  $\sigma' \in BR$  such that  $\sigma'_{\leq m+1} = (\sigma_1, \dots, \sigma_m, t)$ .

Since  $M, \sigma, m \models C_i(n)$ , we have that agent  $i$  has performed exactly  $n$  copy actions on the prefix  $(\sigma_1, \dots, \sigma_m)$ . Moreover, the action that agent  $i$  performs between  $\sigma_m$  and  $t$  is to read  $\alpha$  from  $K_i$ , therefore it has still performed exactly  $n$  copy actions on the prefix  $(\sigma_1, \dots, \sigma_m, t)$ . Then, it is straightforward that  $M, \sigma', m+1 \models C_i(n)$ . That gives us  $M, \sigma', m+1 \models \bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge B^r \alpha \wedge C(i, n)$ . Since  $\sigma'_{\leq m} = \sigma_{\leq m}$  and  $R(\sigma_m, t)$ , we obtain  $M, \sigma, m \models EX(\bigwedge_{\alpha_q \in Q} B_i \alpha_q \wedge C_i(n) \wedge B^r \alpha)$ .

To prove completeness, a satisfying model for a consistent formula is constructed as in the completeness proof of *PCTL\** from [27]. Then we use the axioms to show that this model is in *BMCM*  $(K_1, \dots, K_{n_{AG}}, n_M, n_C)$ .  $\square$

#### 4.4 Systems of heterogeneous reasoners

Changing the logic to accommodate agents which reason using a different set of inference rules rather than resolution is relatively straightforward. As an illustration, we show how to add model conditions and axioms for agents which use modus ponens and conjunction introduction. We assume that the knowledge base of these reasoners contains literals and implications of the form  $L_1 \wedge \dots \wedge L_n \rightarrow L$ .

First of all, we need to change the conditions on models so that instead of using the *Res* action, a agent could change its state by performing *MP* and *AND* actions. Let  $i$  be an (*MP*, *AND*) reasoner. Define  $\Omega_i$  as  $K_i$  closed under subformulas and the following conjunction introduction: if  $Q$  is a set of distinct literals from  $K_i$ , then  $\bigwedge Q \in \Omega_i$ . An agent  $i$  has actions  $Read_{i,\phi,\psi}$  for any formula  $\phi$  in  $K_i$  and  $\psi \in \Omega$  (where  $\psi$  is the overwritten formula),  $Copy_{i,\phi,\psi}$  for any formula  $\phi \in \Omega_i$ ,  $Null_{i,\phi}$ ,  $Erase_{i,\phi}$ , and instead of *Res* it has  $MP_{i,\phi_1,\phi_1 \rightarrow \phi_2,\psi}$  and  $AND_{i,\phi_1,\phi_2,\psi}$ .

DEFINITION 4 (AVAILABILITY OF *MP* AND *AND*). For any  $s \in S$ :

1.  $MP_{i,\phi_1,\phi_1 \rightarrow \phi_2,\psi} \in R_i(s)$  iff  $\phi_1, \phi_1 \rightarrow \phi_2 \in V^r(s, i) \cup V^c(s, i)$  and  $\psi \in \Omega_i$ .
2.  $AND_{i,\phi_1,\phi_2,\psi} \in R_i(s)$  iff  $\phi_1, \phi_2 \in V^r(s, i) \cup V^c(s, i)$  and  $\psi \in \Omega_i$ .

DEFINITION 5 (EFFECTS OF *MP* AND *AND*). For every  $s \in S$ , the result of performing action  $a$  is defined if  $a \in R_i(s)$  and has the following effect on the resulting state  $t$ :

1. if  $a$  is  $MP_{i,\phi,\phi \rightarrow \phi_2,\psi}$  then  $V^r(t, i) = V^r(s, i) \cup \{\phi_2\} \setminus \{\psi\}$  and  $V^c(s, i) = V^c(t, i) \setminus \{\psi\}$ .
2. if  $a$  is  $AND_{i,\phi_1,\phi_2,\psi}$  iff  $V^r(t, i) = V^c(s, i) \cup \{\phi_1 \wedge \phi_2\} \setminus \{\psi\}$  and  $V^c(s, i) = V^c(t, i) \setminus \{\psi\}$ .

The corresponding axioms for the (*MP*, *AND*) reasoner are as follows:

**A13**  $\bigwedge_{q \in Q} B_i \phi_q \wedge C_i(n) \wedge B_i \phi_1 \wedge B_i(\phi_1 \rightarrow \phi_2) \rightarrow EX(\bigwedge_{q \in Q} B_i \phi_q \wedge C_i(n) \wedge B_i \phi_2)$  where  $Q \subseteq \Omega_i$  with  $|Q| < n_M(i)$

**A14**  $\bigwedge_{q \in Q} B_i \phi_q \wedge C_i(n) \wedge B_i \phi_1 \wedge B_i \phi_2 \rightarrow EX(\bigwedge_{q \in Q} B_i \phi_q \wedge C_i(n) \wedge B_i(\phi_1 \wedge \phi_2))$  where  $Q \subseteq \Omega_i$  with  $|Q| < n_M(i)$ , and  $\phi_1, \phi_2 \in \Omega_i$ .

**A15**  $EX(B_i(\phi_1 \wedge \phi_2) \wedge C_i(n)) \rightarrow (B_i(\phi_1 \wedge \phi_2) \vee (\neg B_i(\phi_1 \wedge \phi_2) \wedge B_i \phi_1 \wedge B_i \phi_2 \wedge C_i(n)))$

**A16**  $EX(B_i \phi_2 \wedge C_i(n)) \rightarrow (B_i \phi_2 \vee (\neg B_i \phi_2 \wedge C_i(n) \wedge \bigvee_{\phi_1 \rightarrow \phi_2 \in K_i} (B_i \phi_1 \wedge B_i(\phi_1 \rightarrow \phi_2))))$  for all  $\phi_2 \notin K_i$ .

Now we can add the conditions and axioms for the (*MP*, *AND*) reasoner to the system for resolution reasoners and obtain an axiomatisation for the heterogeneous system of reasoners.

## 5. VERIFYING RESOURCE BOUNDS

The logic *BMCL* allows us to express precisely how beliefs of a set of resource-bounded agents change over time, and, given a memory and communication bound for each agent, to verify formulas which state that a certain belief will or will not be acquired within a certain number of steps. For example, given a system of two agents with premises  $K_1 = \{\{p_1, p_2\}, \{\neg p_1, p_2\}\}$  and  $K_2 = \{\{p_1, \neg p_2\}, \{p_1, \neg p_2\}\}$ , with bounds  $n_M(1) = 2$ ,  $n_M(2) = 2$  (both agents have 2 memory cells) and  $n_C(1) = 0$ ,  $n_C(2) = 1$  (agent 1 cannot copy anything and agent 2 can copy one clause), we can prove that  $start \rightarrow EX^5 B_2(\{\})$  (i.e., from the start state, the agents can derive the empty clause in 5 steps).

However, rather than deriving such properties by hand, it is more convenient to use an automatic method to verify them. In this section, we describe how the models in *BMCM*  $(K_1, \dots, K_{n_{AG}}, n_M, n_C)$  can be encoded as an input to a model-checker to allow the automatic verification of the properties expressing resource bounds.

### 5.1 Model checker encoding

It is straightforward to encode a *BMCM* model of such a system for a standard model checker, and to verify resource bounds using existing model checking techniques. For the examples reported here, we have used the Mocha model checker [10].

States of the *BMCM* models correspond to an assignment of values to state variables in the model-checker. The state variables representing an agent's memory are organised as a collection of 'cells', each holding at most one clause. For an agent  $i$  with memory bound  $n_M(i)$ , there are  $n_M(i)$  cells. Each cell is represented by a pair of bitvectors, each of length  $k = |PROP|$ , representing the positive and negative literals in the clause in some standard order (e.g., lexicographic order). For example, if *PROP* contains the propositional variables  $A_1, A_2$  and  $A_3$  with index positions 0, 1 and 2 respectively, the clause  $A_1 \vee \neg A_3$  would be represented by two bitvectors: "100" for the positive literals and "001" for the negative literals. This gives reasonably compact states.

Actions by each agent such as reading a premise, resolution and communication with other agents are represented by Mocha *atoms* which describe the initial condition and transition relation for a group of related state variables. Reading a premise ( $Read_{i,\alpha,\beta}$ ) simply sets the bitvectors representing an arbitrary cell in agent  $i$ 's memory to the appropriate values for the clause  $\alpha$ . Resolution ( $Res_{i,\alpha_1,\alpha_2,L,\beta}$ ) is implemented using simple bit operations on cells containing values representing  $\alpha_1$  and  $\alpha_2$ , with the results

# agents	Distrib.	Memory	Comm.	Time
1	Symmetric	3	–	8
2	Symmetric	2, 2	1, 0	6
2	Symmetric	3, 3	1, 0	6
2	Symmetric	3, 3	0, 0	8
2	Symmetric	2, 1	1, 1	9
2	Asymmetric	2, 2	2, 1	7
2	Asymmetric	3, 3	2, 1	7
2	Asymmetric	3, 1	1, 0	8

**Table 1: Tradeoffs between resource bounds**

being assigned to an arbitrary cell in agent  $i$ 's memory. Communication ( $Copy_{i,\alpha,\beta}$ ) is implemented by copying the values representing  $\alpha$  from a cell of agent  $j$  to an arbitrary cell of agent  $i$ . To express the communication bound, we use a counter for each agent which is incremented each time a copy action is performed by the agent. After the counter for agent  $i$  reaches  $n_C(i)$ , the  $Copy_{i,\alpha,\beta}$  action is disabled.

Mocha supports hierarchical modelling through composition of *modules*. A module is a collection of atoms and a specification of which of the state variables updated by those atoms are visible from outside the module. In our encoding, each agent is represented by a module. A particular distributed reasoning system is then simply a parallel composition of the appropriate agent modules.

The specification language of Mocha is *ATL*, which includes *CTL*. We can express properties such as ‘agent  $i$  may derive belief  $\phi$  in  $n$  steps’ as  $EF\ tr(B_i\alpha)$  where  $tr(B_i\alpha)$  is a suitable encoding of the fact that a clause  $\alpha$  is present in the agent’s memory (either as a disjunction of possible values of cell bitvectors, or as a special boolean variable which becomes true when one of the cells contains a particular value, for example all 0s for the empty clause). To obtain the actual derivation we can verify the negation of a formula, for example  $AG \neg tr(B_i\alpha)$ , and use the counterexample trace to show how the system reaches the state where  $\alpha$  is proved.

## 5.2 Examples

Consider a single agent (agent 1) whose knowledge base contains all clauses of the form  $\sim A_1 \vee \sim A_2$  where  $\sim A_i$  is either  $A_i$  or  $\neg A_i$ , and which has the goal of deriving the empty clause. We can express the property that agent 1 will derive the empty clause at some point in the future as  $EF B_1 \{\}$ .

Using the model checker, we can show that deriving the empty clause requires a memory bound of 3 and 8 time steps (see Figure 2).<sup>5</sup> We can also show that these space and time bounds are minimal for a single agent; i.e., increasing the space bound does not result in a shorter proof.

With two agents and a symmetric problem distribution (i.e., each agent has all the premises  $\sim A_1 \vee \sim A_2$ ), we can show that a memory bound of 2 (i.e., the minimum required for resolution) and a communication bound of 1 gives a proof of 6 steps (see Figure 3). Reducing the communication bound to 0 results in no proof, as, with a memory bound of 2 for each agent, at least one clause must be communicated from one agent to the other. Increasing the space bound to 3 (for each agent) does not shorten the proof, though it does allow the communication bound to be reduced to 0 at the cost of increasing the proof length to 8 (i.e., the single agent case). Reducing the total space bound to 3 (i.e., 2 for one agent and 1 for the other, equivalent to the single agent case) increases the number of steps required to find a proof to 9 and the communication bound to 1 for each agent. In effect, one agent functions as a cache for a

<sup>5</sup>The space required for problems of this form is known to be logarithmic in the number of premises [16].

clause required later in the proof, and this clause must be copied in both directions.

If the problem distribution is asymmetric, e.g., if one agent has premises  $A_1 \vee A_2$  and  $\neg A_1 \vee \neg A_2$  and the other has premises  $\neg A_1 \vee A_2$  and  $A_1 \vee \neg A_2$ , then with a memory bound of 2 for each agent, we can show that the time bound is 7, and the communication bound is 2 for the first agent and 1 for the second. Increasing the memory bound for each agent to 3 does not reduce the time bound. However the memory bound can be reduced to 1 and the communication bound reduced to 1 for one agent and 0 for the other, if the time bound is increased to 8 (again this is equivalent to the single agent case, except that one agent copies the clause it lacks from the other rather than reading it). These tradeoffs are summarised in Table 1.

Increasing the size of the problem increases the number of possible tradeoffs, but similar patterns can be seen to the 2-variable case. For example, if the agent’s knowledge base contain all clauses of the form  $\sim A_1 \vee \sim A_2 \vee \sim A_3$ , then a single agent requires a memory bound of 4 and 16 steps to achieve the goal. In comparison, two agents, each with a memory bound of 2, require 13 steps and 4 messages to derive the goal.

While extremely simple, these examples serve to illustrate the interaction between memory, time and communication bounds, and between the resource distribution and the problem distribution.

## 6. RELATED WORK

There exist several approaches to epistemic logic which model reasoners as resource-bounded (not logically omniscient), including the deduction model of belief [24], step logic and active logic [15, 20], algorithmic knowledge [21, 17, 26], and other syntactic epistemic logics [14, 3, 8, 22] where each inference step takes the agent into the next (or some future) moment in time. A logic where the depth of belief reasoning is limited is studied in [19].

A considerable amount of work has also been done in the area of model-checking multi-agent systems (see, e.g., [13, 12]). However, this work lacks a clear connection between the way agent reasoning is modelled in the agent theory (which typically assumes that the agents are logically omniscient) and the formalisations used for model checking, and emphasises correctness rather than the interplay between time, memory and communication, and the ability of agents to derive a certain belief.

The current paper extends the work of [4] which proposed a method of verifying memory and time bounds in a single reasoner which reasons in classical logic using natural deduction rather than resolution. We also extend the work in [7] which analyses a system of communicating rule-based reasoners and verifies time bounds for those systems, but assumes unlimited memory. As far as we are aware, the logic we propose in this paper is the first attempt to analyse time, space and communication bounds of reasoners in one logical system, and verify properties relating to all three resources using a model-checker.

## 7. CONCLUSION

In this paper, we analyse the time, space and communication resources required by a system of reasoning agents to achieve a goal. We give a rigorous definition of the measures for each of those resources, and introduce the epistemic logic *BMCL* in which we can express properties of a system of resource-bounded reasoning agents. In particular, we can express bounds on memory and communication resources as axioms in the logic. We axiomatise a system of agents which reason using resolution (other reasoning systems can be axiomatised in a similar way), prove that the resulting

logic is sound and complete, and show how to express properties of the system of reasoning agents in *BMCL*. Finally, we show how *BMCL* transition systems can be encoded as input to the Mocha model-checker and how properties, such as existence of derivations with given bounds on memory, communication, and the number of inference steps, can be verified automatically.

In future work, we plan to consider logical languages containing primitive operators which would allow us to state the agents' resource limitations as formulas in the language rather than axioms, and consider agents reasoning about each other's resource limitations. We also would like to consider agents reasoning in a simple epistemic or description logic.

**Acknowledgements:** This work was partially supported by EP-SRC project EP/E031226/1.

## 8. REFERENCES

- [1] P. Adjiman, P. Chatalic, F. Goasdoué, M.-C. Rousset, and L. Simon. Distributed reasoning in a peer-to-peer setting. In *Proc. of the 16th European Conference on Artificial Intelligence, ECAI'2004*, pages 945–946. IOS Press, 2004.
- [2] T. Ågotnes and N. Alechina. Knowing minimum/maximum  $n$  formulae. In *Proc. of the 17th European Conference on Artificial Intelligence (ECAI 2006)*, pages 317–321. IOS Press, 2006.
- [3] T. Ågotnes and M. Walicki. Strongly complete axiomatizations of "knowing at most" in standard syntactic assignments. In *Pre-proceedings of the 6<sup>th</sup> International Workshop on Computational Logic in Multi-agent Systems (CLIMA VI)*, 2005.
- [4] A. Albore, N. Alechina, P. Bertoli, C. Ghidini, B. Logan, and L. Serafini. Model-checking memory requirements of resource-bounded reasoners. In *Proc. of the Twenty-First National Conference on Artificial Intelligence (AAAI 2006)*, pages 213–218. AAAI Press, 2006.
- [5] N. Alechina, P. Bertoli, C. Ghidini, M. Jago, B. Logan, and L. Serafini. Verifying space and time requirements for resource-bounded agents. In *Proc. of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2006)*, pages 217–219. IEEE Press, 2006.
- [6] N. Alechina, P. Bertoli, C. Ghidini, M. Jago, B. Logan, and L. Serafini. Verifying space and time requirements for resource-bounded agents. In *Proc. of the Fourth Workshop on Model Checking and Artificial Intelligence (MoChArt-2006)*, pages 16–30, 2006.
- [7] N. Alechina, M. Jago, and B. Logan. Modal logics for communicating rule-based agents. In *Proc. of the 17th European Conference on Artificial Intelligence (ECAI 2006)*, pages 322–326. IOS Press, 2006.
- [8] N. Alechina, B. Logan, and M. Whitsey. A complete and decidable logic for resource-bounded agents. In *Proc. of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 606–613. ACM Press, 2004.
- [9] M. Alekhovich, E. Ben-Sasson, A. A. Razborov, and A. Wigderson. Space complexity in propositional calculus. *SIAM J. of Computing*, 31(4):1184–1211, 2002.
- [10] R. Alur, T. A. Henzinger, F. Y. C. Mang, S. Qadeer, S. K. Rajamani, and S. Tasiran. MOCHA: Modularity in model checking. In *Computer Aided Verification*, pages 521–525, 1998.
- [11] E. Amir and S. A. McIlraith. Partition-based logical reasoning for first-order and propositional theories. *Artificial Intelligence*, 162(1-2):49–88, 2005.
- [12] M. Benerecetti, F. Giunchiglia, and L. Serafini. Model checking multiagent systems. *J. Log. Comput.*, 8(3):401–423, 1998.
- [13] R. Bordini, M. Fisher, W. Visser, and M. Wooldridge. State-space reduction techniques in agent verification. In *Proc. of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2004)*, pages 896–903. ACM Press, 2004.
- [14] H. Duc. Reasoning about rational, but not logically omniscient, agents. *J. Log. Comput.*, 5:633–648, 1997.
- [15] J. J. Elgot-Drapkin and D. Perlis. Reasoning situated in time I: Basic concepts. *J. of Experimental and Theoretical Artificial Intelligence*, 2:75–98, 1990.
- [16] J. L. Esteban and J. Torán. Space bounds for resolution. In *Proc. STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 551–560. Springer, 1999.
- [17] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, Mass., 1995.
- [18] B. Faltings and M. Yokoo. Introduction: Special issue on distributed constraint satisfaction. *Artif. Intell.*, 161(1-2):1–5, 2005.
- [19] M. Fisher and C. Ghidini. Programming Resource-Bounded Deliberative Agents. In *Proc. of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI'99)*, pages 200–206. Morgan Kaufmann, 1999.
- [20] J. Grant, S. Kraus, and D. Perlis. A logic for characterizing multiple bounded agents. *Autonomous Agents and Multi-Agent Systems*, 3(4):351–387, 2000.
- [21] J. Y. Halpern, Y. Moses, and M. Y. Vardi. Algorithmic knowledge. In *Proc. of the 5th Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 255–266. Morgan Kaufmann, 1994.
- [22] M. Jago. *Logics for Resource-Bounded Agents*. PhD thesis, University of Nottingham, 2006.
- [23] H. Jung and M. Tambe. On communication in solving distributed constraint satisfaction problems. In *Multi-Agent Systems and Applications IV, Proc. 4th International Central and Eastern European Conference on Multi-Agent Systems, CEEMAS 2005*, pages 418–429. Springer, 2005.
- [24] K. Konolige. *A Deduction Model of Belief*. Morgan Kaufmann, San Francisco, Calif., 1986.
- [25] G. M. Provan. A model-based diagnosis framework for distributed embedded systems. In *Proc. of the Eighth International Conference on Principles and Knowledge Representation and Reasoning (KR-02)*, pages 341–352. Morgan Kaufmann, 2002.
- [26] R. Pucella. Deductive algorithmic knowledge. In *AI&M 1-2004, Eighth International Symposium on Artificial Intelligence and Mathematics*, 2004.
- [27] M. Reynolds. An axiomatization of *PCTL\**. *Inf. Comput.*, 201(1):72–119, 2005.
- [28] M. Wooldridge and P. E. Dunne. On the computational complexity of coalitional resource games. *Artif. Intell.*, 170(10):835–871, 2006.
- [29] A. C.-C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Conference Record of the Eleventh Annual ACM Symposium on Theory of Computing*, pages 209–213. ACM, 1979.